FINAL REPORT

Tensor Invariant Processing for Munitions/Clutter Classification

SERDP Project MR-2100

NOVEMBER 2013

Thomas Bell
Science Applications International Corporation

This document has been cleared for public release

SERDP
DOD ▪ EPA ▪ DOE

Report Documentation Page	Form Approved OMB No. 0704-0188

1. REPORT DATE **NOV 2013**	2. REPORT TYPE	3. DATES COVERED **00-00-2013 to 00-00-2013**
4. TITLE AND SUBTITLE **Tensor Invariant Processing for Munitions/Clutter Classification**		5a. CONTRACT NUMBER
		5b. GRANT NUMBER
		5c. PROGRAM ELEMENT NUMBER
6. AUTHOR(S)		5d. PROJECT NUMBER
		5e. TASK NUMBER
		5f. WORK UNIT NUMBER
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) **Science Applications International Corporation,1710 SAIC Drive ,McLean,VA,22102**		8. PERFORMING ORGANIZATION REPORT NUMBER
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)		10. SPONSOR/MONITOR'S ACRONYM(S)
		11. SPONSOR/MONITOR'S REPORT NUMBER(S)
12. DISTRIBUTION/AVAILABILITY STATEMENT **Approved for public release; distribution unlimited**		
13. SUPPLEMENTARY NOTES		

14. ABSTRACT

Objectives The intent of this research project is to explore alternatives to conventional dipole inversion for extracting target features from multi-axis EMI sensor array data. Conventional dipole inversion searches for the target parameter values (location and polarizabilities) which minimize the difference between measured signals and those calculated using the dipole response model. Here, we consider an alternative approach that seeks to determine the parameter values which minimize an objective function based on the dispersion in estimates of the target?s polarizability using different combinations of transmitters and receivers. Technical Approach The alternative inversion approach exploits rotationally invariant properties of the polarizability tensor. Focusing the processing on the primary invariant (trace of the polarizability tensor) allows us to significantly limit the search space that is required to invert EMI data. Conventional dipole inversion requires that we search over target (x, y, z) location, target (θ, ϕ, ψ) orientation and (β_1, β_2, β_3) principal axis polarizabilities to minimize the difference between the measured response and response predicted by the dipole model fit. With tensor invariant processing we only need to search on (x, y, z) to find the target location which minimizes the dispersion in calculated values for the rotationally invariant trace. The principal axis polarizabilities can then be calculated directly. Results This report documents results on the convergence properties of a downhill simplex based algorithm for determining a target?s location (and polarizabilities) using this approach. Specifically we examined convergence of the algorithm for data collected with the 2x2 TEM array at the former Camp Beale in 2011 and found no impact due to signal-to-noise ratio (SNR) and background leveling effects. However, the minimum polarizability dispersion does vary systematically with SNR ? targets with higher SNR tend to have less uncertainty (dispersion) in the polarizability estimate than those with lower SNR. Testing shows that polarizability dispersion based inversion can produce more accurate polarizabilities than conventional dipole inversion in some cases. However, the fraction of targets in the ESTCP classification demonstrations that cannot be accurately analyzed using conventional dipole inversion may be small enough (<5%) that modest improvements in calculating the polarizability are likely to have very little effect on classification performance. We also document a new size-shape classification algorithm for comparing unknown target polarizabilities with those of munitions items and other targets of interest. It appears to produce better classification performance than the library comparison procedure that we have been using. We also compare the performance of the size-shape classification

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE	**Same as Report (SAR)**	**28**	
unclassified	**unclassified**	**unclassified**			

Abstract

Objectives

The intent of this research project is to explore alternatives to conventional dipole inversion for extracting target features from multi-axis EMI sensor array data. Conventional dipole inversion searches for the target parameter values (location and polarizabilities) which minimize the difference between measured signals and those calculated using the dipole response model. Here, we consider an alternative approach that seeks to determine the parameter values which minimize an objective function based on the dispersion in estimates of the target's polarizability using different combinations of transmitters and receivers.

Technical Approach

The alternative inversion approach exploits rotationally invariant properties of the polarizability tensor. Focusing the processing on the primary invariant (trace of the polarizability tensor) allows us to significantly limit the search space that is required to invert EMI data. Conventional dipole inversion requires that we search over target (x, y, z) location, target (θ, ϕ, ψ) orientation and $(\beta_1, \beta_2, \beta_3)$ principal axis polarizabilities to minimize the difference between the measured response and response predicted by the dipole model fit. With tensor invariant processing we only need to search on (x, y, z) to find the target location which minimizes the dispersion in calculated values for the rotationally invariant trace. The principal axis polarizabilities can then be calculated directly.

Results

This report documents results on the convergence properties of a downhill simplex based algorithm for determining a target's location (and polarizabilities) using this approach. Specifically we examined convergence of the algorithm for data collected with the 2x2 TEM array at the former Camp Beale in 2011 and found no impact due to signal-to-noise ratio (SNR) and background leveling effects. However, the minimum polarizability dispersion does vary systematically with SNR – targets with higher SNR tend to have less uncertainty (dispersion) in the polarizability estimate than those with lower SNR.

Testing shows that polarizability dispersion based inversion can produce more accurate polarizabilities than conventional dipole inversion in some cases. However, the fraction of targets in the ESTCP classification demonstrations that cannot be accurately analyzed using conventional dipole inversion may be small enough (<5%) that modest improvements in calculating the polarizability are likely to have very little effect on classification performance.

We also document a new size-shape classification algorithm for comparing unknown target polarizabilities with those of munitions items and other targets of interest. It appears to produce

better classification performance than the library comparison procedure that we have been using. We also compare the performance of the size-shape classification algorithm using the average polarizability (equivalent to the trace of the polarizability tensor) with the performance for the full set of three principal axis polarizabilities using data from the ESTCP demonstrations at Camp San Luis Obispo, Camp Butner, Camp Beale and Pole Mountain. The differences are not large, but the 3β ROC tends to rise a bit faster than the <β> ROC. Because the two approaches can emphasize slightly different features in the EMI response, different targets tend to drive the ROC behavior of the different classifiers.

Benefits

By its very nature, polarizability dispersion based inversion provides a direct measure of the uncertainty in the polarizabilities calculated from EMI data collected using advanced sensor arrays. The benefits of the new classification algorithm are both quantitative and qualitative. Re-processing data from the recent Camp Beale demonstration using this approach produced a Receiver Operating Characteristic (ROC) which rises more rapidly and hits the 100% TOI recovered level with 50% fewer clutter digs beyond the training set than the ROC from conventional processing. Improved classification performance improves munitions response efficiency. The procedure operates in an intuitive and easily visualized feature space. It is transparent, objective and easily automated. All of this is likely to facilitate transition to production work and ease regulatory acceptance.

Contents

Figures

Acronyms

EMI	Electromagnetic Induction
ESTCP	Environmental Security Technology Certification Program
FOM	Figure of Merit
ISO	Industry Standard Object
ROC	Receiver Operating Characteristic
Rx	Receive
SERDP	Strategic Environmental Research and Development Program
SNR	Signal to Noise Ratio
TEM	Transient Electromagnetic
TOI	Target of Interest
Tx	Transmit

Objective

The Strategic Environmental Research and Development Program (SERDP) and the Environmental Security Technology Certification Program (ESTCP) have invested heavily in developing advanced electromagnetic induction (EMI) sensor technology capable of providing the information needed to reliably distinguish between buried munitions items and metallic clutter in the ground such as munitions fragments, scraps of exploded metal objects, cultural or agricultural artifacts and debris, etc. The processing and analysis techniques that are used to extract the information needed to reliably classify targets as munitions or clutter from data collected with these sensors were originally developed for use with an older, less sophisticated generation of EMI sensors such as the Geonics EM61. The newer sensors provide a richer, multi-dimensional view of the target which is not fully exploited by the processing schemes that are currently used.

The objectives of this project are to

  a) Improve understanding of the relationships between multi-axis EMI sensor array data and the intrinsic features of target response,
  b) Develop new processing approaches for multi-axis EMI sensor data that make full use of the capabilities available with newly developed EMI sensor technology, and
  c) Improve procedures for identifying targets as munitions or clutter.

Our basic approach involves exploring new processing techniques that can be used to calculate the standard target features used for classification (principal axis polarizabilities), but that do not rely on conventional dipole inversion. In conventional dipole inversion, given a set of measurements of the EMI response we simultaneously search out the location, orientation, and principal axis polarizabilities that produce signals that best match the measured target response. With the new multi-axis technology, the transmit (Tx) and receive (Rx) fields can be combined in simple ways to selectively excite and observe the target's response in a set of orthogonal directions. This can be accomplished using appropriate linear combinations of the data recorded in the sensor's various transmit and receive channels, and allows the polarizabilities to be calculated directly once the target's location has been determined. Different combinations of transmit and receive coils correspond to different views of the target, and we need only find the target location for which those different views produce consistent results for the target parameters.

This report documents results on the convergence properties of our algorithm for determining a target's location (and polarizabilities) using this approach. Specifically we examine the effects of signal-to-noise ratio (SNR) and background leveling on convergence using data collected at the former Camp Beale in 2011. We also document a size-shape classification algorithm for comparing unknown target polarizabilities with those of munitions items and other targets of

interest, which appears to produce better classification performance than the library comparison procedure that we have been using.

Technical Approach

Electromagnetic induction is the process whereby changing magnetic fields create currents in electrically conducting objects. Transient electromagnetic (TEM) sensors create a magnetic field by running a current through a coil. When the current is shut off this primary field shuts down, inducing currents in any nearby metallic objects. The induced currents quickly decay as they diffuse through an object. TEM sensors measure the voltage induced in a coil by the decaying magnetic field associated with these eddy currents.

Dipole Response Model

In the standard dipole response model, we represent the target response by an induced dipole moment $\mathbf{m}(t)$ which is proportional to the primary field $\mathbf{H}_0$ at the target location. The proportionality factor is the target's magnetic polarizability tensor $\mathbf{B}(t)$. The elements B_{ij} of $\mathbf{B}$ form a 3×3 array of dipole moment components in the i coordinate directions induced by excitation in the j coordinate directions such that

$$m_x(t) = B_{xx}(t)H_{0x} + B_{xy}(t)H_{0y} + B_{xz}(t)H_{0z},$$

$$m_y(t) = B_{yx}(t)H_{0x} + B_{yy}(t)H_{0y} + B_{yz}(t)H_{0z},$$

$$m_z(t) = B_{zx}(t)H_{0x} + B_{zy}(t)H_{0y} + B_{zz}(t)H_{0z}.$$

Receive coils measure the signal from the induced dipole field. The sensitivity to the induced dipole components mimics the field that would be produced by currents flowing through the receive coil. Both the primary field and the receive coil response are calculated as Biot-Savart integrals around the coils. Writing $\mathbf{T}$ for the field of the transmit coil and $\mathbf{R}$ for the pseudofield of the receive coil (including receiver gain factors), the signal from the transmit/receive coil pair is then simply

$$S = \mathbf{R} \cdot \mathbf{B}\mathbf{T}.$$

$\mathbf{B}$ is a second rank tensor that depends on the size, shape and material properties of the target, as well as its orientation relative to the x, y, z coordinate directions. The target orientation relative to the x, y, z coordinate axes is specified by a set of roll, pitch and yaw angles ψ, θ and ϕ. If the coordinate system is rotated by the angles ψ, θ and ϕ into alignment with the principal axes of the target, then the off diagonal elements of $\mathbf{B}$ go to to zero ($B_{ij} = 0$ for $i \neq j$), and the diagonal elements correspond to the eigenvalues β_i of $\mathbf{B}$. These are referred to as the target's principal axis polarizabilities.

Using the dipole response model we can calculate the polarizability from EMI data collected over a target [1]. The target is classified as munitions (or other target of interest) or clutter depending on whether its polarizability is more munitions-like or clutter-like. Conventional dipole inversion uses an iterative search procedure to determine the dipole response model parameters which produce the closest match between the model and the measured response. The process is illustrated as a flow diagram in Figure 1 below.
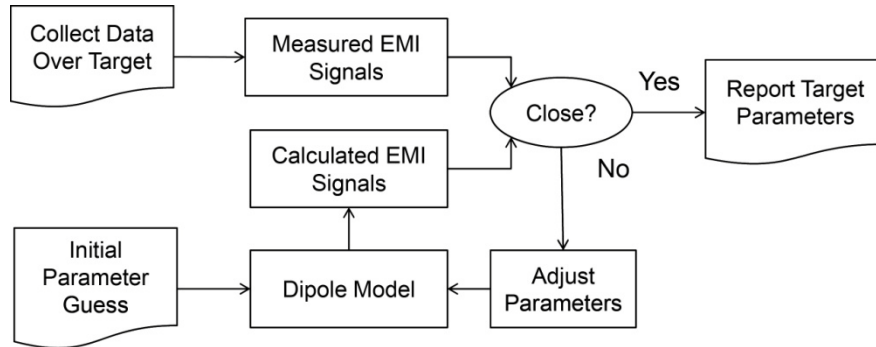

Figure 1. Flow diagram for conventional dipole inversion.

This procedure was originally developed for data collected at various locations over a target with a monostatic (a single pair of co-located transmit and receive coils) EMI sensor such as the Geonics EM61. The new generation of EMI sensors purpose-built for classification are fixed arrays of transmit and receive coils which excite and observe the target over a broad range of angles. This is illustrated in two dimensions for a pair of different coil configurations in Figure 2. The green and blue lines are field lines for a pair of coils A and B. The pairs of arrows at the target show the directions of the primary field vectors $\mathbf{H}_A$ and $\mathbf{H}_B$ from the two coils.


Figure 2. Multi-coil arrays exciting target with different primary field directions.

In general, the fields from three coils arranged in some reasonable fashion can be combined to create field components of unit strength in three orthogonal directions at the target location. The coefficients of the linear combinations are solutions of the linear vector equation

$$\mathbf{a} \otimes \mathbf{H}_A + \mathbf{b} \otimes \mathbf{H}_B + \mathbf{c} \otimes \mathbf{H}_C = \mathbf{I}$$

where $\mathbf{I}$ is the identity matrix and the symbol $\otimes$ signifies the outer product of the coefficient and field vectors. This is actually three sets of equations wherein linear combinations of fields $\mathbf{H}$ from coils A, B and C with coefficients a, b and c produce unit vectors $\mathbf{n}_x$, $\mathbf{n}_y$ and $\mathbf{n}_z$ in x, y and z directions at the target, e.g.

$$a_x H_{Ax} + b_x H_{Bx} + c_x H_{Cx} = 1$$

$$a_x H_{Ay} + b_x H_{By} + c_x H_{Cy} = 0$$

$$a_x H_{Az} + b_x H_{Bz} + c_x H_{Cz} = 0$$

for the x direction. A well-conditioned H field component matrix

$$\mathbf{H} = \begin{bmatrix} H_{Ax} & H_{Bx} & H_{Cx} \\ H_{Ay} & H_{By} & H_{Cy} \\ H_{Az} & H_{Cz} & H_{Cz} \end{bmatrix}$$

implies good coil/position combination for target illumination or observation.

The signal from a single transmit/receive coil pair is $S = \mathbf{R} \cdot \mathbf{BT}$. Given a set of three transmit and three receive coils, we can solve for coefficients p_{ij} and q_{ij} for transmit and receive coil combinations which synthesize orthonormal transmit field and receive pseudo-field components at the target location and then directly calculate the components of the polarizability tensor

$$B_{ij} = \sum_{k,l} p_{ik} S(Tx_k, Rx_l) q_{jl}.$$

This presents an alternative to conventional dipole inversion. So long as we use the correct target location we should get the same polarizability components for any transmit/receive combination. The idea then is to search out the target location which minimizes the dispersion of polarizability estimates from different sets of transmit and receive coils. Significantly, the minimum dispersion provides a direct measure of the uncertainty in our polarizability estimate.

Minimizing Polarizability Dispersion

As an expedient, we use the primary invariant of the polarizability tensor for our calculations. There are three rotationally invariant quantities associated with the polarizability tensor [2]. The trace Tr($\mathbf{B}$) is equal to the sum of the diagonal elements, and because it is a rotational invariant, also equal to the sum of the eigenvalues. The other two invariants are the determinant of $\mathbf{B}$ and the sum of the determinants of the 2x2 diagonal sub-matrices of $\mathbf{B}$. By virtue of their rotational invariance, these are equal respectively to the product of the eigenvalues and the sum of products

of the eigenvalues taken pairwise. The invariants are the coefficients of the characteristic polynomial of **B**, and the principal axis polarizabilities can be directly calculated from them if desired. An average polarizability <β> can be calculated by averaging the polarizability tensor **B** over all orientation angles. The average polarizability is a scalar function of time that describes the overall average response of the target. It is equal to the average of the principal axis polarizabilities, and hence equal to ⅓ times the trace of **B** [3].

The basic algorithm is illustrated schematically in Figure 3. We iteratively search out the target location which minimizes the dispersion in estimates of <β> for a variety of coil combinations. Once the target location has been determined we can go on to calculate the other invariants or the full polarizability tensor if desired.
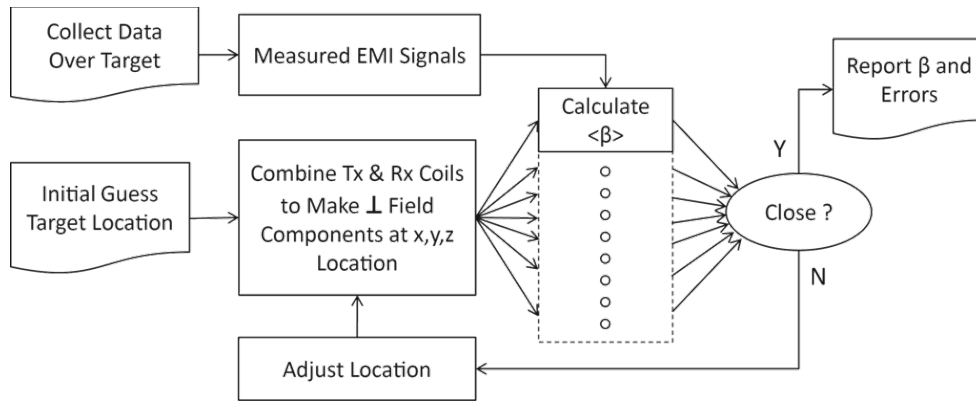

Figure 3. Flow chart for alternative approach to estimating the polarizability of a target.

There are N!/[3! x (N-3)!] ways of combining N transmit or receive coils in sets of three. The 2x2 TEM array [4] has four transmit coils which yield four Tx sets and four three-axis receiver cubes (12 coils in all) which yield 220 Rx sets. Thus with this array we could calculate the polarizability 880 different ways. Not all of the possible ways are equally viable, and there is clearly redundancy since the 2x2 array produces only 48 independent measurements of the EMI response.

Our processing algorithm selects a subset based on the signal strength and how well the receivers interrogate the nominal target location. These are determined by the norms and the condition numbers of the field component matrices of the different possible combinations. The norm $||\mathbf{H}||$ of the field component matrix **H** is a gauge of the field strength at the target. We use the Euclidian (L_2) norm, given by the largest eigenvalue of **H**. It is proportional to the loop area and the number of turns, and more or less inversely proportional to the third power of the range to the target. The condition number $k = ||\mathbf{H}|| \, ||\mathbf{H}^{-1}||$, which is equal to the ratio of the largest to the smallest eigenvalue, is a gauge of coil combination quality for target illumination/observation. When $k = 1$ all principal axes are being interrogated equally well. Large values of k indicate that at least one of the axes is not being interrogated very well. Larger $||\mathbf{H}||$ and smaller k make for

better coil combinations. Figure 4 shows the distribution of field matrix norms and condition numbers for the 2x2 receiver array with a shallow target near the center of the array.


Target x = -1.5 cm, y = 3 cm, d = 5 cm

Figure 4. 2x2 receiver array field matrix norm and condition number for a shallow target near the center of the array.

We use a subset of the 880 possible transmit/receive sets (typically ~25-50) with the largest norms and smallest condition numbers for the nominal target location in our calculations – i.e. those to the lower right in Figure 4. The error surface for target location using our approach is appears to be well behaved. Figure 5 shows the polarizability dispersion error surface over the x-y plane at the target depth for target BE-529 (a 60mm mortar) from the 2011 Camp Beale demonstration [5]. Here we use one set of three transmit coils and 25 sets of receive coils, and calculate the dispersion of the net polarizability $P = \sum_{i,j} B_{ii}(t_j)$ as the ratio of the standard deviation of the P values for the 25 receiver combinations to the mean, $\sigma(P)/\overline{P}$. In this example, we sum over all time gates $t_j > 88 \ \mu s$. The minimum of the error surface provides a direct measure of the uncertainty in the polarizability estimate.


BE-529

Figure 5. Polarizability dispersion error surface for target location for Camp Beale target BE-529.

7

We use the basic Nelder and Mead downhill simplex algorithm [6] to search for the target location which minimizes the polarizability dispersion. Once the target location is determined we calculate the principal axis polarizabilities. Figure 6 compares principal axis polarizabilities calculated by minimizing polarizability dispersion (lines) with polarizabilities from standard dipole inversion (points) for two of the Camp Beale targets. BE-407 is a horseshoe and BE-529 is a 60 mm mortar. For BE-529, the dipole inversion calculation took 4.02 s using IDL on a Lenovo Thinkpad X201, while the polarizability dispersion minimization took 3.12 s. The effects of signal-to-noise ratio (SNR) and background leveling on convergence are analyzed using data from the Camp Beale demonstration.


Figure 6. Principal axis polarizabilities calculated by minimizing polarizability dispersion (lines) compared with polarizabilities from standard dipole inversion (points)

Signal to Noise and Background Leveling

With the 2x2 TEM array, noise and background leveling are intertwined. At early times the background response is dominated by coil ring-down effects, while at later times electromagnetic noise effects become more important. Background (no target) signals are subtracted from data collected over a target before processing. When the sensor is being used in the static mode for cued data collection, standard practice is to take background shots every 20-30 minutes. Variability of the background signals provides an estimate of the noise level. We have chosen a subset of the 2x2 TEM data from the Camp Beale demonstration for analysis. Figure 7 shows the site. The boxes show areas where 2x2 TEM data were collected, and the dots show the locations which appeared to be target-free and were used for background shots. Our analysis uses data collected on June 8, 2011 in the open area outlined in red with green dots. There were 14 background shots spread out over a 6-hour period. Figure 8 shows the mean background level for the various transmit/receive combinations. The monostatic (z axis receivers along the block diagonal) combinations show the most pronounced ringdown. Standard deviations of the background signals for the various transmit/receive combinations are shown in Figure 9. The

8

apparent noise for the monostatic combinations is a bit different from the others, possibly reflecting some contribution from soil response variations over the region.


Figure 7. Ariel photograph of Camp Beale site showing 2x2 TEM data collection areas and background shot locations.

Data were collected over 152 targets during this time period. Figure 10 shows the distribution of signal to noise ratio (SNR) vs. decay time for these targets. The median and interquartile ranges are shown by the red curves. SNR is calculated as the average over all of the 48 signal channels (four transmitters and four three-axis receivers) of the background subtracted signal divided by the corresponding noise estimate from Figure 9. Note that the SNR peaks at about 0.1 ms.

We tested convergence of the new algorithm on these data using as a metric the number of function calls to calculate the polarizability dispersion from the downhill simplex algorithm. The results are given later in the Results and Discussion section.
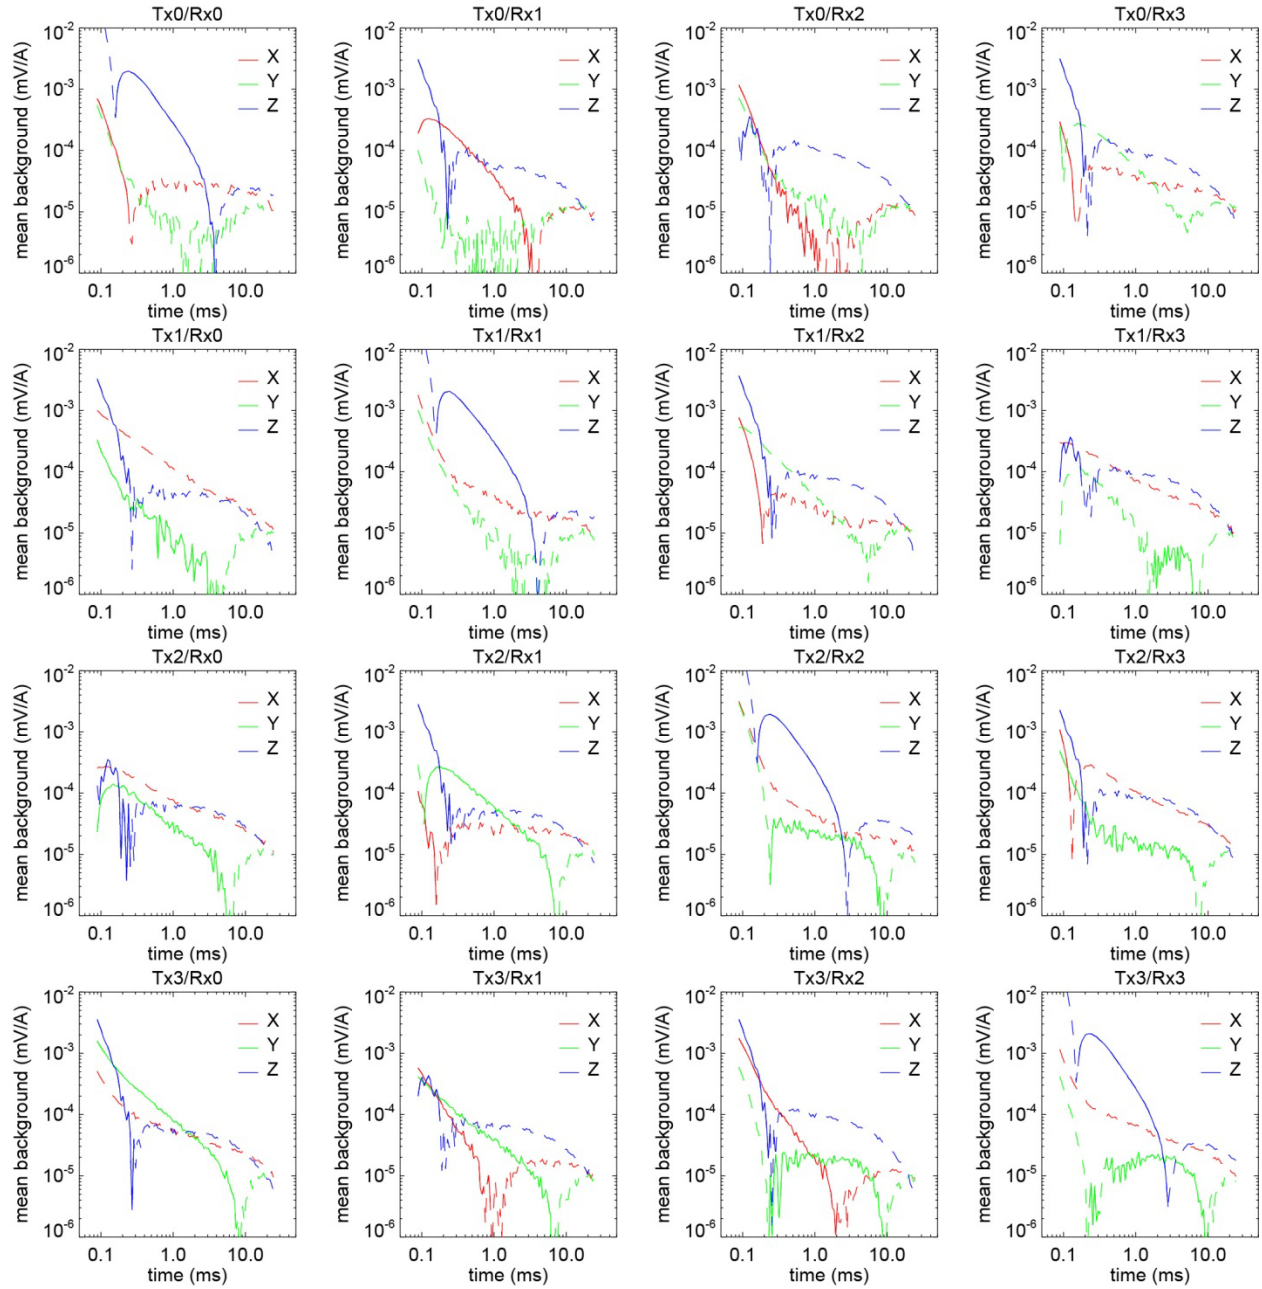
Figure 8. Mean background levels for 2x2 TEM data collected on June 8, 2011.

Figure 9. Standard deviation of background levels for 2x2 TEM data collected on June 8, 2011

Figure 10. Signal to noise ratio vs. decay time for 2x2 TEM target data collected at Camp Beale on June 8, 2011.

Classification

Classification is a matter of deciding whether the object's polarizabilities are munitions-like or clutter-like. Library matching methods employing various procedures to compare polarizabilities of unknown targets with those of targets of interest (TOIs) are commonly used for classification [7, 8]. Ours exploits the fact that an object's polarizability tensor $\mathbf{B}(t)$ can be represented as a product of two factors: the volume V of the object and a tensor $\mathbf{A}(t)$ whose eigenvalues $\alpha_j(t)$, $i = 1, 2, 3$ depend only on the shape and composition of the object. Strictly speaking this is true only for nonmagnetic objects [9] or those with a specific magnetic permeability. However, as a practical matter it appears to be a good representation for typical TOIs and clutter items. Confronted with an unknown target, we then compare its apparent size and EMI "shape" with the sizes and shapes of the TOI.

Given the set (spanning three axes and N time gates) of principal axis polarizabilities β_{ref} for a TOI and the set of principal axis polarizabilities β for an unknown target, we calculate a size ratio s as

$$s = \text{median}\left(\frac{\sqrt[3]{\beta}}{\sqrt[3]{\beta_{ref}}}\right).$$

Here β and β_{ref} are 3N element arrays, and calculations are performed element-by-element. The median is taken over all axes and time gates for which $\beta > 0$. Since β scales with volume, $\beta^{1/3}$ scales with the linear dimension of the target, i.e. its size. With noise, standard inversion procedures can produce negative polarizabilities at some time gates. If $\beta < 0$ for more than some threshold fraction (typically 25-50%) of the available terms, then the target is put in the "can't analyze" category. The size mismatch parameter Δ_{size} is then

$$\Delta_{size} = \log(s)$$

which is equal to zero if the EMI sizes of the target and the reference TOI are the same, The shape mismatch parameter Δ_{shape} is determined by comparing the unknown target's polarizability with the reference polarizability scaled by the size mismatch

$$\Delta_{shape} = \frac{\sum\left|\sqrt[3]{\beta} - s\sqrt[3]{\beta_{ref}}\right|}{\sum \sqrt[3]{\beta}}$$

in which the sums are over all terms with positive β. We retain the cube root to increase the influence of the later time gates, for which the response is substantially weaker than it is at the earlier time gates. For each target, size and shape mismatch parameters are calculated for each TOI. Classification is based on applying a threshold to a figure of merit (FOM) parameter

$$\text{FOM} = \min_{\text{TOI}}\left\{|\Delta_{size}| + k\log(\Delta_{shape})\right\}.$$

13

We find that using $k \approx 0.3$ seems to give the best overall classification performance. Minimizing the FOM over the set of TOI then finds the best match to any TOI.

Figure 11 shows the results of applying this procedure to polarizabilities from the Camp Beale demonstration of the 2x2 TEM array. Because the focus now is on comparing our classification algorithm with that used for the original dig list submission [5], we use polarizabilities calculated with standard dipole inversion during the demonstration. The top plot shows the distribution of size and shape mismatch parameters for the targets in the final SAIC dig list. Plotted values are for the minimum FOM over the TOI set. The TOI are small pipe sections called "Industry Standard Objects" (ISOs) which were used as seed items, 37 mm projectiles, 60 and 81 mm mortars (two varieties each) and 105 mm projectiles. The TOI signatures used for classification were determined using a clustering algorithm based on size and shape mismatches between all pairs of targets in the data set. The symbols are color coded retrospectively using target IDs from the scoring report. In this feature space the targets of interest cluster nicely in the region near $(\Delta_{size}, \Delta_{shape}) = (0, 0)$, while the clutter items do not. The receiver operating characteristic (ROC) curve for classification performance using our new procedure (blue curve) is compared with the ROC for the final SAIC dig list (red-orange-green curve) in the bottom plot. Although the "traditional" processing represented by our dig list submission did quite well, the new approach seems to do somewhat better. The new ROC curve rises more rapidly and hits the 100% TOI recovered level with fewer clutter digs beyond the training set than the original ROC. The blue curve does not highlight the "can't analyze" targets (there are 7 vs. 32 with the "traditional" processing), nor does it indicate a stop-dig threshold. In practice the stop-dig threshold will be determined using training data picks spread over a "halo" near the origin in our size/shape mismatch feature space. The halo will be determined empirically from the FOM distribution combined with polarizability error estimates based on the polarizability dispersion discussed in the previous section.

It is convenient to work with the average polarizability $< \beta >= \frac{1}{3}Tr(\mathbf{B})$ when inverting EMI data by minimizing the uncertainty in the polarizability estimate. We have also considered whether or not it is useful for classification by comparing classification performance using <β> with the performance using all three βs. The calculations use the same formulae, but we simply apply the algorithm using the N element array <β> rather than the 3N element β array. Our results are presented and discussed in the next section.
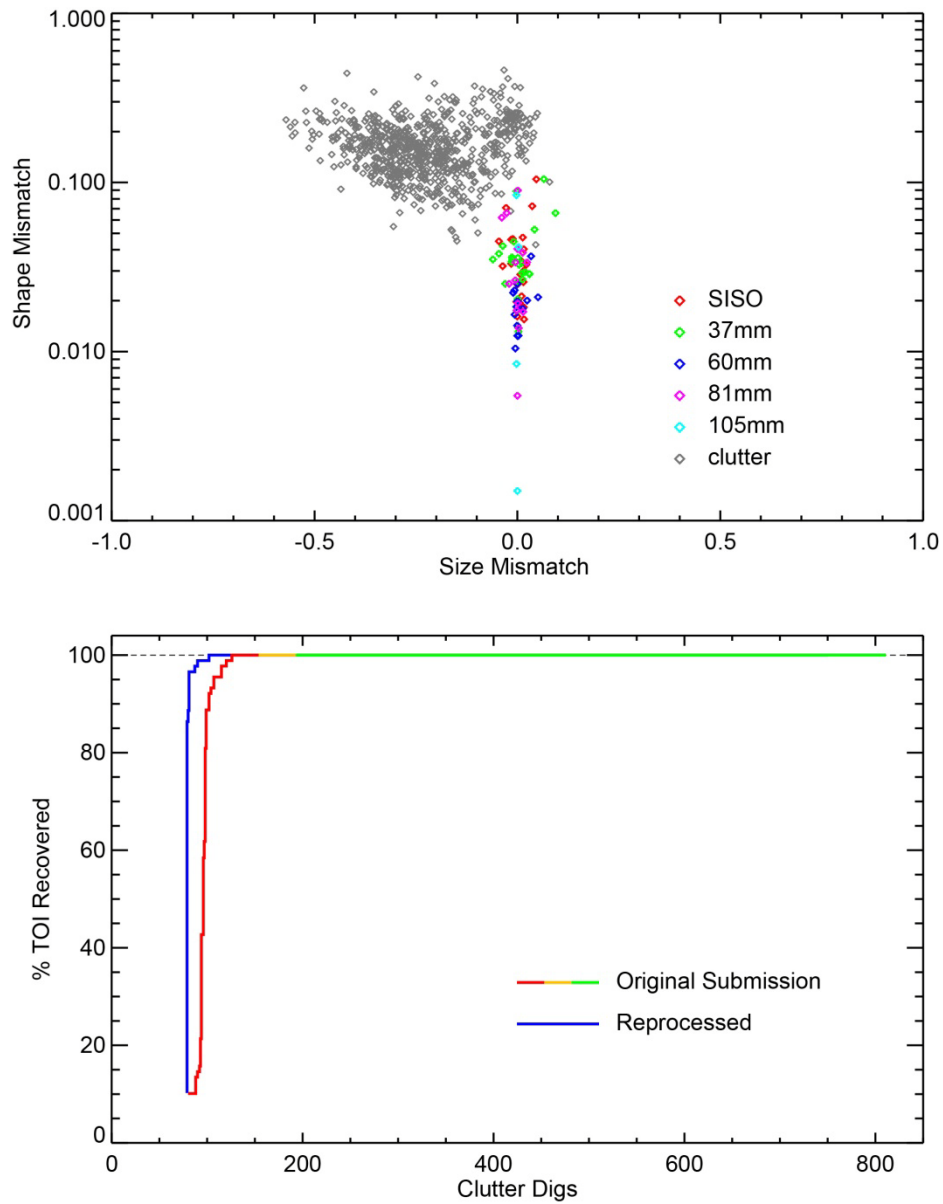
Figure 11. Top: Minimum FOM size and shape mismatch parameters for the targets in the final SAIC Camp Beale dig list. Bottom: ROC curves for the original classification using "traditional" processing (red-orange-green) and re-classification using our new approach (blue).

15

Results and Discussion

Convergence of the new algorithm was tested using single target anomalies from the Camp Beale data. This is appropriate since as currently configured the algorithm assumes a single target. Like the standard dipole inversion, if there is more than one target it will normally still converge to an answer which represents some hybrid of the targets' responses. Again, like the standard dipole inversion it could be generalized to multi-target anomalies, but it is not capable of directly determining how many targets are contributing to the EMI response. Using only those anomalies for which only one target was found during the post-test intrusive investigation leaves 106 of the original 152 anomalies.

The downhill simplex algorithm needs a fairly good starting point, so we tested using initial guesses for the target location under each of the four quadrants of the array, at x = ±15 cm and y = ±15 cm (the centers of the sensors in the array are at x = ±20 cm and y = ±20 cm) and at 15 cm depth. Polarizability dispersion was calculated over all four transmit combinations and twelve receive combinations, averaged over time gates ranging from 0.1-0.6 ms. Using the four starts, the algorithm converged for all of the targets. Most of the anomalies (66 of the 106) converged from all of the quadrants, while only two converged from just one quadrant. Of the remaining anomalies, 31 only converged from three of the four quadrants and seven converged from only two. The number of function calls to calculate the polarizability dispersion is a gauge of how rapidly the algorithm converges, and is plotted as a function of SNR for the 106 single-target anomalies in Figure 12. All four quadrant starts are included, with a limit of 150 calls each. If a target converged for none of the quadrant starts there would be 600 calls. The median number of calls per target is 382, and there is no obvious trend with SNR.


Figure 12. Function calls to calculate polarizability vs. signal to noise ratio.

The minimum polarizability dispersion $\sigma(P)/\overline{P}$ does vary with SNR (Figure 13, left). SNR is calculated as the average over all of the 48 signal channels of the background subtracted signal

16

voltage divided by the corresponding noise estimate from Figure 9. As one might expect, targets with higher SNR tend to have less uncertainty (dispersion) in the polarizability estimate than those with lower SNR. Although there is a good bit of spread, the polarizability dispersion based inversion data seems to be reasonably well fit by a straight line in log-log space of the form $\sigma(P)/\overline{P} \propto SNR^{-1}$, consistent with the Cramer-Rao bound on parameter estimation error [10].


Figure 13. Residual error vs. signal to noise ratio for polarizability dispersion and standard dipole inversion.

Fit error for conventional dipole inversion (basically $\sigma(data - model)/\overline{data}$) generally follows a similar trend [11], but there is a bit more scatter in the plot of dipole fit error vs. SNR for these data (Figure 13, right) than in the corresponding polarizability dispersion plot on the left. There is little correlation (r = 0.22) between the minimum polarizability dispersion and the mismatch between data and model with conventional dipole inversion. This most likely reflects the fact that the two approaches treat the data differently, but the reason for the increased scatter with the dipole fit error remains unclear.

With conventional dipole inversion the standard metric for data-model match is the mean squared deviation, in which case it is a separable nonlinear least squares problem [12]. The signal S depends linearly on the polarizability **B**, but the transmit field **T** and receive pseudofield **R** used to calculate the signal both depend nonlinearly on the target location relative to the sensor. The usual approach with such problems is to iteratively search on the nonlinear parameters (target x, y, z location in this case) while solving the linear least squares problem for the linear parameters (the elements of the polarizability tensor) at each step. With the 2x2 TEM array there is an overdetermined set of 48 linear equations in six unknowns (**B** is symmetric so only six of its nine elements are unique) to be solved at each step. The solution minimizes the mean squared error in matching all of the data. With polarizability dispersion based inversion we use a similar tactic of iteratively searching for the target location which minimizes an objective function, but now that objective function reflects the uncertainty in the polarizability estimate rather than the actual mismatch between observed and modeled signals. The polarizability

17

calculations are equivalent to solving 6-equation subsets of the 48 which use three different transmitters and three different receivers. We are clearly weighting the data differently with the two inversion approaches, and should expect somewhat different behavior.

In at least some cases we see improved performance using the polarizability based objective function. Figure 14 shows results of test stand measurements of a 37 mm projectile with the 2x2 TEM array. The plot shows the errors in calculated polarizabilities for the two processing techniques for increasing vertical separation between the sensors and the projectile. The polarizability error plotted in this figure is the absolute difference between the calculated $\beta_i(t)$ and reference 37 mm $\beta_i(t)$ at 62 time gates ranging from 0.09 ms to 2.77 ms, normalized by the reference polarizability. For the reference βs we use the average of the standard dipole inversion βs and the polarizability based βs at 29 cm. Beyond 50 cm we see an obvious increase in the β error for standard dipole inversion (black diamonds) relative to our polarizability based processing (blue circles), although both are increasing rather rapidly at this point. The improvement is due to the fact that the polarizability based processing is getting a better location estimate. But the advantage is soon lost. At greater depths the target is simply not being illuminated well enough by the sensor.


Figure 14. Errors in calculating the polarizability of a 37 mm projectile at various distances below the 2x2 TEM array with standard dipole inversion and polarizability dispersion based processing.

The performance of the size-shape classification algorithm relative to the performance of standard library matching was discussed in the Technical Approach section. We have also compared the performance of the size-shape classification algorithm for average polarizability (equivalent to the trace of the polarizability tensor) with the performance for the full set of three principal axis polarizabilities using data from the ESTCP demonstrations at Camp San Luis Obispo, Camp Butner, Camp Beale and Pole Mountain (year 1). The Camp San Luis Obispo and Camp Butner data used here were collected with the vehicle-towed 5x5 TEM array [13, 14], the

Camp Beale data with the 2x2 manportable TEM array [5] and the Pole Mountain data with the MetalMapper multiaxis TEM sensor array [15].

The results are shown as ROC curves in Figure 15. The blue curves show 3β classification performance and the red curves show <β> performance. The differences are not large, but the 3β ROC tends to rise a bit faster than the <β> ROC. Different targets drive the ROC behavior of
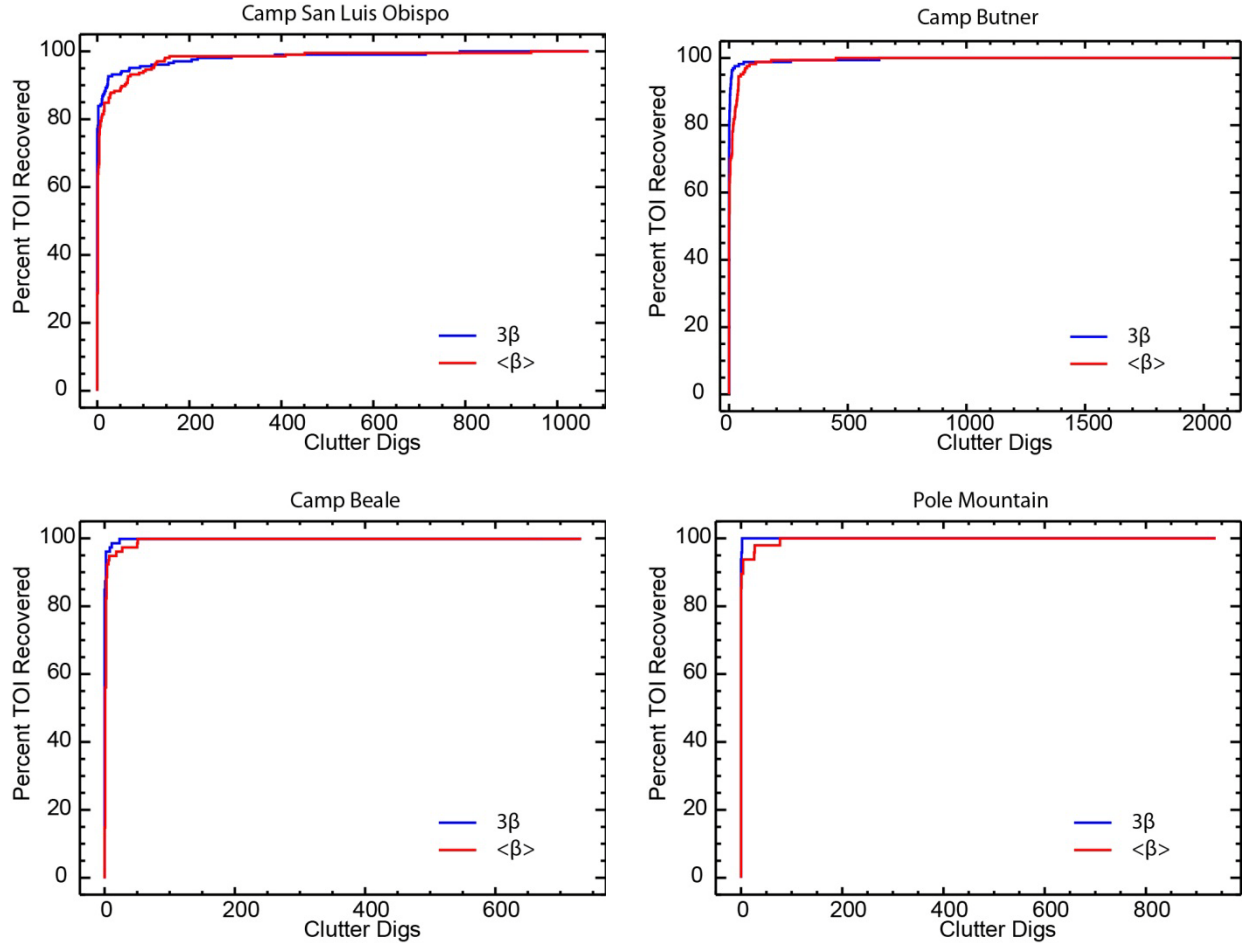

Figure 15. ROC curves for 3β and <β> classification at different ESTCP demonstration sites.

the two classifiers. Figure 16 is a comparison of the Camp Beale target FOM values for classifiers based on all three polarizability eigenvalues and on the average polarizability. Blue circles correspond to munitions items and other targets of interest and red triangles correspond to clutter items. The plot on the right shows the detail in the transition region between TOI and clutter. The vertical dashed line passes through the "most difficult" TOI for the 3β classifier, while the horizontal dashed line passes through a different TOI, which is "most difficult" for the <β> classifier. The clutter items in the lower right quadrant come before the last TOI is captured with the <β> classifier, but after the last TOI is captured with the 3β classifier. Those in the upper left quadrant come before the last TOI is captured with the 3β classifier, but after the last TOI is captured with the <β> classifier. The 3β rises more rapidly because there are fewer clutter

items in the upper left than in the lower right. Much of the latter are exploded ordnance fragments that are a better match to the small seeded ISO pipe sections in the <β> classifier than in the 3β classifier.


Figure 16. Comparison of target FOM values for classifiers based on all three polarizability eigenvalues and on the average polarizability. Detail on right. Blue circles correspond to munitions items and other targets of interest and red triangles correspond to clutter items.

Conclusions

The intent of this research project was to explore alternatives to conventional dipole inversion for extracting target features from multi-axis EMI sensor array data. Conventional dipole inversion searches for the target parameter values (location and polarizabilities) which minimize the difference between measured signals and those calculated using the dipole response model. Here, we have considered an alternative approach that seeks to determine the parameter values which minimize an objective function based on the dispersion in estimates of the target's polarizability using different combinations of transmitters and receivers.

This report documents results on the convergence properties of a downhill simplex based algorithm for determining a target's location (and polarizabilities) using this approach. Specifically we examined convergence of the algorithm for data collected at the former Camp Beale in 2011 and found no impact due to signal-to-noise ratio (SNR) and background leveling effects. The algorithm converges properly for all field targets irrespective SNR. However, the minimum polarizability dispersion does vary systematically with SNR – targets with higher SNR tend to have less uncertainty (dispersion) in the polarizability estimate than those with lower SNR.

Testing shows that polarizability dispersion based inversion can produce more accurate polarizabilities than conventional dipole inversion in some cases. However, the fraction of targets in the ESTCP classification demonstrations that cannot be accurately analyzed using conventional dipole inversion may be small enough (<5%) that modest improvements in calculating the polarizability are likely to have very little effect on classification performance. In the next phase of this project we will be exploring improved procedures for target classification based on principal axis polarizabilities.

Finally, this report describes a size-shape classification algorithm for comparing unknown target polarizabilities with those of munitions items and other targets of interest. It appears to produce better classification performance than the library comparison procedure that we have been using. We also compare the performance of the size-shape classification algorithm using the average polarizability (equivalent to the trace of the polarizability tensor) with the performance for the full set of three principal axis polarizabilities using data from the ESTCP demonstrations at Camp San Luis Obispo, Camp Butner, Camp Beale and Pole Mountain. The differences are not large, but the 3β ROC tends to rise a bit faster than the $<\beta>$ ROC. Because the two approaches can emphasize slightly different features in the EMI response, different targets tend to drive the ROC behavior of the different classifiers.

Literature Cited

1. Thomas H. Bell, Bruce J. Barrow and Jonathan T. Miller, "Subsurface discrimination using electromagnetic induction sensors," *IEEE Trans. Geoscience and Remote Sensing*, vol. 39, no. 6, 1286-1293, June 2001.

2. Jacob Korevaar, *Mathematical Methods*, Vol. 1, §10.2, Academic Press, New York, 1968.

3. J.F. Douglas and B.E. Garboczi, "Intrinsic viscosity and the polarizability of particles having a wide range of shapes," *Adv. Chem. Phys.*, 91, 85-153, 1995.

4. James B. Kingdon, Bruce J. Barrow, Thomas H. Bell, David C. George, Glenn R. Harbaugh and Daniel A. Steinhurst, "TEMTADS Adjunct Sensor Systems - Hand-held EMI Sensor for Cued UXO Discrimination (ESTCP MR-200807) and Man-Portable EMI Array for UXO Detection and Discrimination (ESTCP MR-200909) Final Report," U.S. Naval Research Laboratory Report NRL/MR/6110--12-9401, April 2012.

5. ESTCP Live Site Demonstrations, Former Camp Beale, Marysville, CA, Demonstration Plan Supplement, TEMTADS 2x2 Cart Survey, March 2011.

6. J.A. Nelder and R. Mead, "A simplex method for function minimization," *Computer Journal*, vol. 7, no. 4, 308-313, 1965

7. Leonard R. Pasion, Stephen D. Billings, Douglas W. Oldenburg and Sean E. Walker, "Application of a library based method to time domain electromagnetic data for the identification of unexploded ordnance," Journal of Applied Geophysics, vol. 61 pp. 279–291, 2007.

8. Karl N. Kappler, "An approach to UXO discrimination via polarizability curvematching and feature extraction applied to polarizability," 21st Symposium on the Application of Geophysics to Environmental and Engineering Problems (SAGEEP), Philadelphia, April 2008.

9. L.D. Landau, E.M. Lifshitz and L.P. Pitaevskii, Electrodynamics of Continuous Media, 2nd Edition, §59, Elsevier, Amsterdam, 1984.

10. Harry L. van Trees, *Detection, Estimation and Modulation Theory*, Part 1, §2.7, John Wiley and Sons, New York, 1968.

11. Final Report, SERDP Project MM-1310, "Sensor Orientation Effects on UXO Geophysical Target Discrimination," December 2006.

12. Gene Golub and Victor Pereyra, "Separable nonlinear least squares: the variable projection method and its applications," *Inverse Problems*, vol. 19, no. 2, R1-R26, 2003.

13. D.A. Steinhurst, G.R. Harbaugh, J.B. Kingdon, T. Furuya, D.A. Keiswetter and D.C. George, "EMI Array for Cued UXO Discrimination," Final Report, ESTCP Project MM-0601, July 2010.

14. Nagi Khadr, James B. Kingdon, Glenn R. Harbaugh and Daniel A. Steinhurst, "Demonstration Data Report, Former Camp Butner, MTADS Discrimination Array (TEMTADS) Survey," Naval Research Laboratory Report NRL/MR/6110--11-9366, October 20, 2011.

15. Stephen Billings, "Data Collection with Vehicular-Based Systems - Pole Mountain, WY," Demonstration Report, ESTCP Project MR-201160, September 2012.